

Computing Like the Brain: The path to machine intelligence

NASA

September 17, 2013

Jeff Hawkins

jhawkins@numenta.com



1) Discover operating principles of neocortex

2) Build systems based on these principles

Artificial Intelligence - no neuroscience

Alan Turing



“Computers are universal machines”

1935+

“Human behavior as test for machine intelligence”

1950

Major AI Initiatives

- MIT AI Lab
- 5th Generation Computing Project
- DARPA Strategic Computing Initiative
- DARPA Grand Challenge



AI Projects

- ACT-R
- Asimo
- CoJACK
- Cyc
- Deep Blue
- Global Workspace Theory
- Mycin
- SHRDLU
- Soar
- Watson
- Many more -




Pros: - Good solutions




Cons: - Task specific
- Limited or no learning

Artificial Neural Networks – minimal neuroscience

Warren McCulloch
Walter Pitts





“Neurons as logic gates” 1943
Proposed first artificial neural network



ANN techniques

- Back propagation
- Boltzman machines
- Hopfield networks
- Kohonen networks
- Parallel Distributed Processing
- Machine learning
- Deep Learning



Pros:

- Good classifiers
- Learning systems


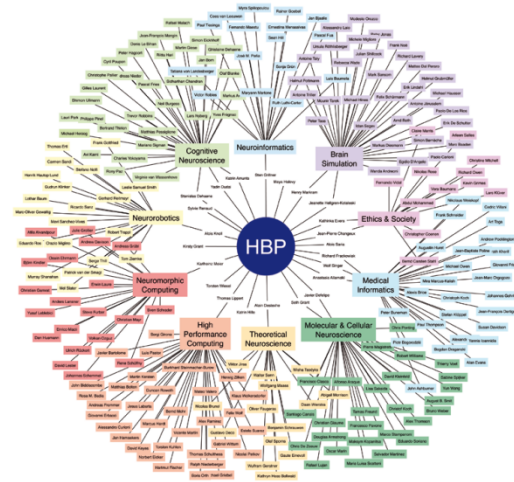

Cons:

- Limited capabilities
- Not brain like

Whole Brain Simulator – maximal neuroscience

The Human Brain Project

Spatial scales




Blue Brain simulation

No theory


No attempt at Machine Intelligence

1) Discover operating principles of neocortex

2) Build systems based on these principles




The neocortex is a memory system.




The neocortex learns a sensory-motor model of the world

Principles of Neocortical Function

1) On-line learning from streaming data




Principles of Neocortical Function




1) On-line learning from streaming data

2) Hierarchy of memory regions


Principles of Neocortical Function



Principles of Neocortical Function




Principles of Neocortical Function




- 1) On-line learning from streaming data
- 2) Hierarchy of memory regions
- 3) Sequence memory
- 4) Sparse Distributed Representations
- 5) All regions are sensory and motor

Principles of Neocortical Function



- 1) On-line learning from streaming data
- 2) Hierarchy of memory regions
- 3) Sequence memory
- 4) Sparse Distributed Representations
- 5) All regions are sensory and motor
- 6) Attention

Principles of Neocortical Function



- 1) On-line learning from streaming data
- 2) Hierarchy of memory regions
- 3) Sequence memory
- 4) Sparse Distributed Representations
- 5) All regions are sensory and motor
- 6) Attention

These six principles are necessary and sufficient for biological and machine intelligence.

- All mammals from mouse to human have them
- We can build machines like this

- $$01101101 = m$$

- [illegible]

SDR Properties

- 1) **Similarity:**
shared bits = semantic similarity

[illegible]

- 2) **Store and Compare:**
store indices of active bits

[illegible]

subsampling is OK

Indices

- ### 3) Union membership:

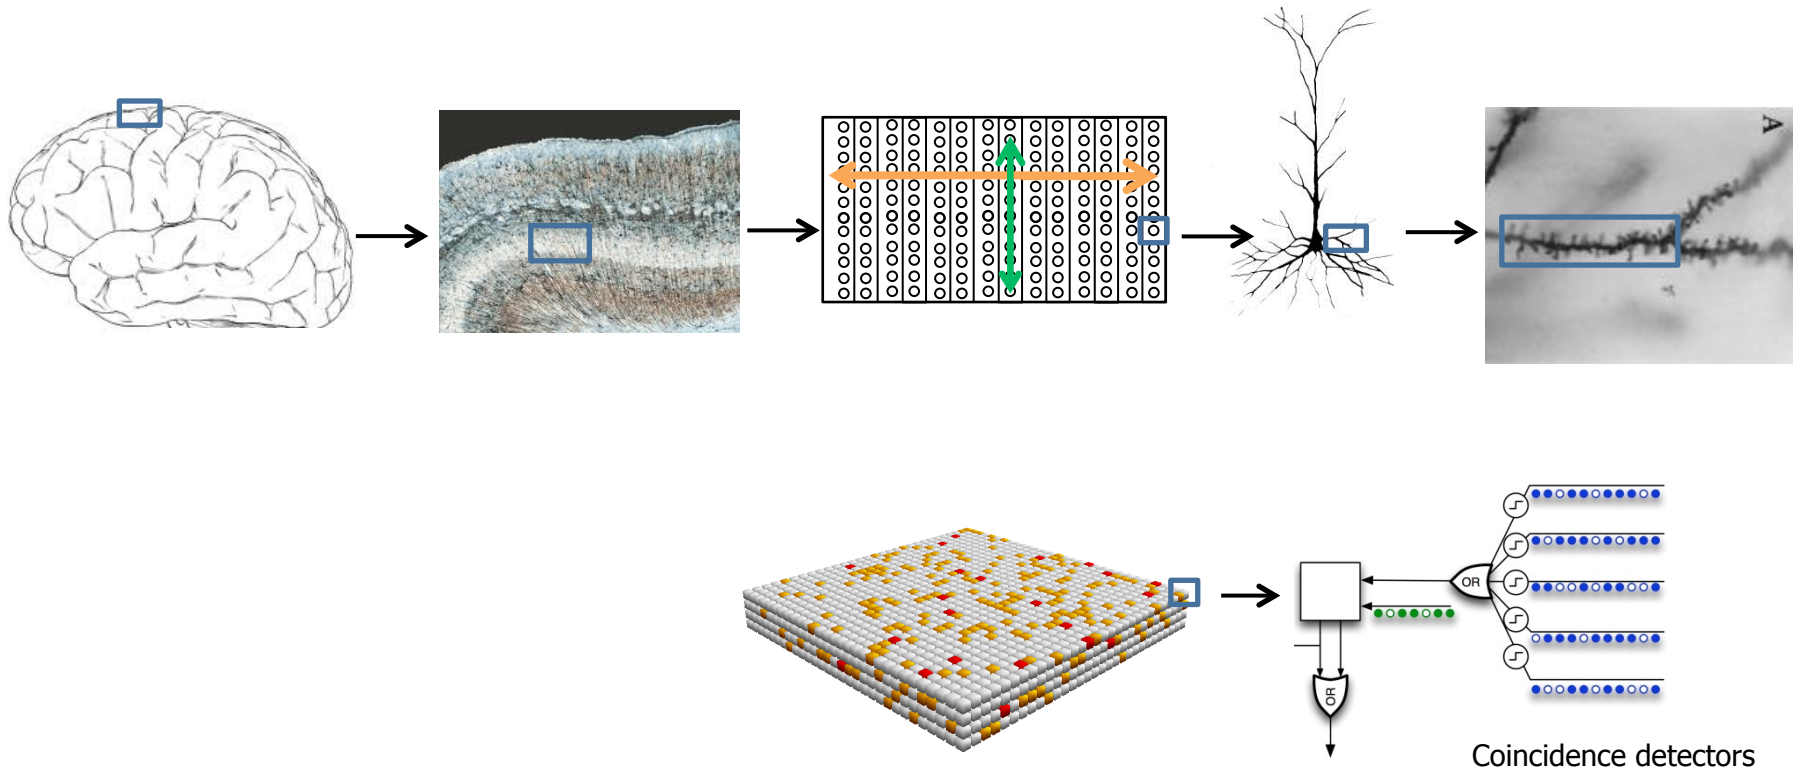
[illegible]

Union

Is this SDR
a member?


[illegible]

Sequence Memory (for inference and motor)



How does a layer of neurons learn sequences?


Each cell is one bit in our Sparse Distributed Representation




SDRs are formed via a local competition between cells.

All processes are local across large sheets of cells.



SDR (time =1)




SDR (time =2)



Cells form connections to subsample of previously active cells.
Predicts its own future activity.




Multiple Predictions Can Occur at Once





With one cell per column, 1st order memory
We need a high order memory

High order sequences are enabled with multiple cells per column.



High Order Sequence Memory




40 active columns, 10 cells per column

= 10^{40} ways to represent the same input in different contexts

A-B-C-D-E

X-B'-C'-D'-Y

High Order Sequence Memory




Distributed sequence memory
Works across large areas
High order, high capacity
Multiple simultaneous predictions
Semantic generalization

Online learning

- Learn continuously, no batch processing
- If pattern repeats, reinforce, otherwise forget it



Learning is the growth of new synapses.





Connection strength is binary



Connection permanence is a scalar

Training changes permanence


Cortical Region



Cortical Region



Cortical Region



2 mm

	sequence memory	CLA	Feedforward inference
	sequence memory	CLA	Feedforward inference
	sequence memory	CLA	Motor output
	sequence memory	CLA	Feedback / attention

Evidence suggests each layer is implementing a CLA variant

Three Current Directions

1) Open Source Project

- NuPIC: CLA open source software and community
- Improve algorithms, develop applications

2) Commercialization

- GROK: Predictive analytics using CLA
- Commercial value generates investment \$

3) Custom CLA Hardware

- Needed for scaling research and commercial applications
- IBM, Seagate, Sandia Labs, DARPA

NuPIC: CLA Open Source Project


www.Numenta.org

Single source tree (used by GROK)

GPLv3 license


Active community

- 215 mail list subscribers
- 20 messages per day
- growing
- full time manager, Matt Taylor

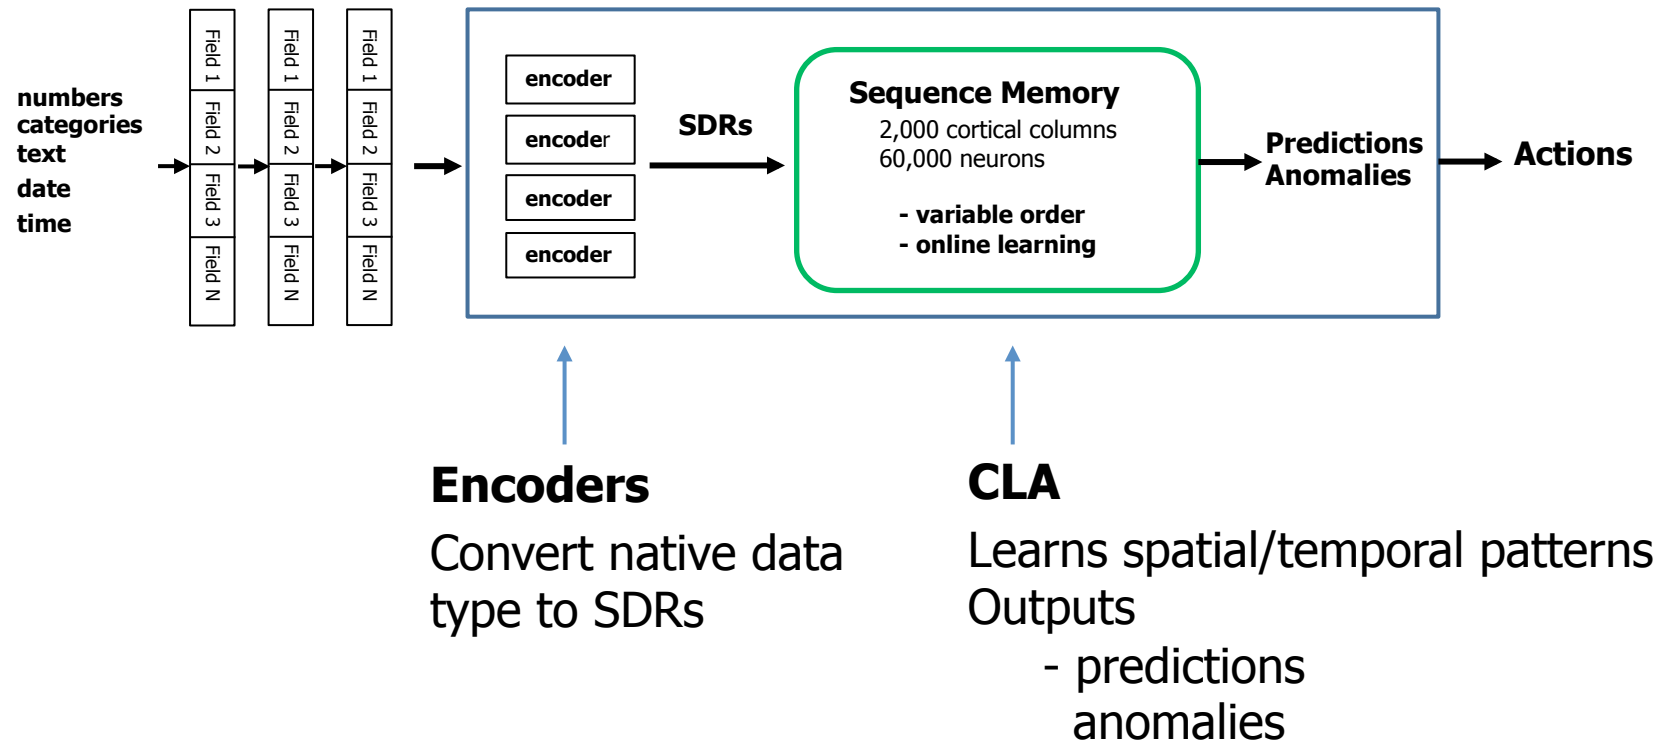


Next hackathon November 2 & 3 in San Francisco


- NLP using SDRs
- Sensory-motor integration using CLA discussion




GROK: Predictive Analytics Using CLA



GROK example: Factory Energy Usage




Customer need



At midnight, make 24
hourly predictions


GROK Predictions and Actuals



GROK example: Predicting Server Demand

Grok used to predict server demand


Approximately 15% reduction in AWS cost



Server demand, Actual vs. Predicted

GROK example: Detecting Anomalous Behavior

Grok builds model of data, detects changes in predictability.



Gear bearing temperature & Grok Anomaly Score

GROK going to market for anomaly detection in I.T. 2014

Custom CLA Hardware

IBM

- Almaden Labs
- Joint research agreement
- Winfried Wilcke


DARPA

- "Cortical Processor"
- "HTM" (Hierarchical Temporal Memory)
- CLA is prototype primitive
- Dan Hammerstrom

Seagate

Sandia Labs

Future of Machine Intelligence



Future of Machine Intelligence



Definite

- **Faster, Bigger**
- **Super senses**
- **Fluid robotics**
- **Distributed hierarchy**



Maybe

- **Humanoid robots**
- **Computer/Brain interfaces for all**



Not


- **Uploaded brains**
- **Evil robots**
- **Friendly uses only**



Why Machine Intelligence?



Live better



Learn more

Thank You